HP ProLiant DL585 Server technology

technology brief, 2nd edition

Abstract

The HP ProLiant DL585 server is a 4U rack-optimized, 4-way server created for large data center deployments requiring enterprise-class performance, uptime, and scalability, plus easy management and expansion. Based on AMD™ Opteron™ technology, the ProLiant DL585 server greatly improves the performance of standard 32-bit applications. Moreover, it includes AMD 64-bit extensions that add the capability to run 32-bit and 64-bit applications on the same x86 platform with no performance penalty.

This paper is intended for IT professionals familiar with system administration and ProLiant industry-standard servers. It explains the technologies implemented in the ProLiant DL585 server that make it an excellent choice for enterprise customers needing increased performance while maintaining investment protection with a 64-bit migration path.

Acronyms in text

The following acronyms are used in the text of this document.

Table 1. Acronyms

Acronym	Acronym expansion
ASIC	Application specific integrated circuit
BBWC	Battery-backed write cache
CPU	Central processing unit
CRM	Customer relationship management
DIMM	Dual inline memory module
EDA	Electronic design automation
ERP	Enterprise resource planning
HPTC	High performance technical computing
iLO	Integrated Lights-Out
IPF	Itanium processor family
LOM	LAN on motherboard
NIC	Network interface card
NFT	Network fault tolerance
NUMA	Non-uniform memory access
PCI	Peripheral component interconnect
PCI-X	Peripheral component interconnect extended
RBSU	ROM-based setup utility
SAS	Serial attached SCSI
SATA	Serial ATA
SFF	Small form factor
SUMA	Sufficiently-uniform memory access

Introduction

The Opteron-based ProLiant DL585 server is a 4-way, x86 rack-optimized server delivering best-in-class performance, industry-leading management, and enterprise-class reliability for large data center deployments. The DL585 is a 4-way platform for high-performance 32-bit and 64-bit computing requirements. This capability allows IT organizations to protect their large x86 investments and manage their 32- to 64-bit transition.

Customers can order the ProLiant DL585 server in a 2P configuration or upgrade it to support four processors. Computing environments for which the server is particularly well suited include:

- High-performance technical computing (HPTC)
- Electronic design automation (EDA)/semiconductor
- Financial
- Large database applications
- ERP/CRM
- Petrochemical
- Life science and material science
- Video rendering

Operating system support

The ProLiant DL585 server supports the following operating systems for single-core processing[1]:

- Microsoft® Windows® 2000 Server
- Microsoft Windows 2000 Advanced Server
- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Red Hat Enterprise Linux 3
- Red Hat Enterprise Linux 4
- SuSE Linux Enterprise Server 8
- SuSE Linux Enterprise Server 9
- VMware Virtualization Software

The ProLiant DL585 server supports the following operating systems for dual-core processing (update or service pack for dual-core processing listed if required)[2]:

- Microsoft® Windows® 2000 Server (licensed by core)
- Microsoft Windows 2000 Advanced Server (licensed by core)
- Windows Server 2003, Standard Edition (licensed by processor socket)
- Windows Server 2003, Enterprise Edition (licensed by processor socket)
- Red Hat Enterprise Linux 3 Update 5
- Red Hat Enterprise Linux 4 Update 2
- SuSE Linux Enterprise Server 9 SP1
- VMware Virtualization Software

[1] For up-to-date information on supported operating systems, visit the HP OS Support matrices at: http://www.hp.com/go/supportos

[2] For an up-to-date listing of the latest drivers available for the HP ProLiant DL585 server, see www.hp.com/support/files

ProLiant DL585 architecture

The ProLiant DL585 server is a powerful, 4U enterprise server incorporating technologies that extend the capabilities of industry-standard x86 computing. This server provides high performance and the capability of running both 32-bit and 64-bit applications simultaneously with no performance penalty when using an operating system that supports 64-bit extensions. This is possible because the x-86-64 instruction set architecture is a super-set containing the x86-32 instruction set architecture.

For the ProLiant DL585 server, HP designed a new system architecture that supports Opteron processors and other advanced technologies that enterprise customers have come to expect from ProLiant servers. The ProLiant DL585 is the first links-based x86 server platform. It provides excellent system performance by using powerful processors, on-chip memory controllers, a large memory footprint, and processor-to-processor communication to significantly reduce latency. Figure 1 is a block diagram of ProLiant DL585 architecture.

Figure 1. ProLiant DL585 architecture

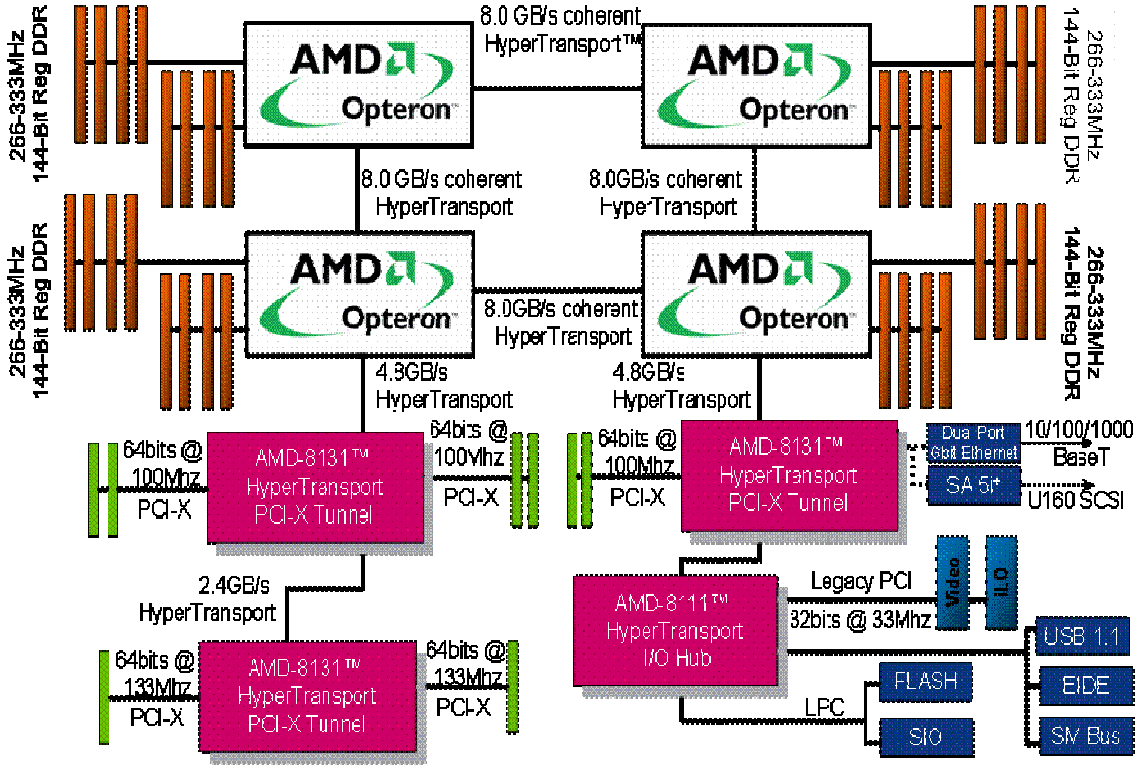

The essential features of the ProLiant DL585 architecture include:

- AMD Opteron 800 series, single-core or dual-core processors and the AMD-8000™ series chipset supporting up to 1 GHz HyperTransport™
- Embedded memory controllers and 144-bit wide memory bus
- AMD-8111 HyperTransport I/O hub (south bridge)
- Eight full-length, 64-bit PCI-X I/O slots

Please refer to the HP ProLiant DL585 server QuickSpecs for additional details, available at this URL: http://h18004.www1.hp.com/products/quickspecs/11902_div/11902_div.html

AMD-8000™ series chipset

The ProLiant DL585 server uses the AMD-8000 series chipset with Direct Connect Architecture and HyperTransport™ technology. HyperTransport is a parallel point-to-point interconnect that replaces parallel front-side bus technology. Direct Connect architecture is AMD's designation for the coherent HyperTransport connection between processors. It eliminates the bottlenecks inherent in front-side bus technology by integrating the memory controller into the AMD Opteron chip and directly connecting CPUs to the IO subsystem and other processors.

The chipset consists of the following key components:

- AMD-8131 HyperTransport PCI-X tunnels, Three are used in the DL585.
- AMD-8111 HyperTransport I/O hub (south bridge)

The AMD Opteron includes the AMD64 instruction set architecture, which enables the ProLiant DL585 server to support both 32-bit and 64-bit applications running simultaneously—with no performance penalty. Unlike the Itanium processor family (IPF) which requires the mode of operation to switch between x86 and IPF modes, the AMD64 instruction set is an extension of the x86 instruction set and requires no mode change to support 64-bit instructions.

Dual-core technology

A dual-core processor is a single die that contains two CPU cores, each with its own 1-MB L2 cache. The dual-core technology leverages the same memory and HyperTransport technology resources available in the single-core processors. Dual-core technology delivers high performance and reduced latency for multithreaded and multitasking environments while maintaining a similar power requirement to single-core processors. Dual-core processors generally run at one or two clock-steps below the current top speed single-core version, thereby maintaining the same power envelope. For more details about AMD processor architecture, refer to the HP white paper "The AMD processor roadmap for HP ProLiant servers," available at this URL:
http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00428708/c00428708.pdf

Processor/memory boards

The ProLiant DL585 server supports up to four processor/memory boards (2 or 4 are the only valid configurations supported). Each of these boards includes an AMD Opteron 800 series processor running at up to 2.4 GHz for dual-core processors or up to 2.8 GHz, for single-core processors. There is an integrated memory controller that operates at CPU core frequency, a 144-bit-wide memory bus, and eight DIMM slots supporting up to 128 GB of memory, assuming 4 GB DIMMs.

The Opteron processors operate at core speeds of up to 2.8 GHz. Each processor supports 64-KB integrated Level 1 (L1) instruction cache, 64-KB integrated L1 data cache, and 1-MB L2 cache. The processors support IA-32 and the AMD-64 instruction set for running 64-bit applications.

Memory subsystem

Typical multiprocessor PC server architecture connects IA-32 processors to memory DIMMs by means of a north bridge chip. The north bridge is a memory controller and bridge to the I/O expansion interface. AMD Opteron architecture differs from typical multiprocessor server architecture in that AMD's north bridge is integrated into the Opteron 800 series processor to boost performance. Performance is enhanced by eliminating bus contention created when memory and I/O paths pass through a typical, non-integrated north bridge. Because the memory controller is integrated onto the processor chip, memory latency is greatly reduced.

In the ProLiant DL585, processors access the memory controllers at core speed. Each on-chip memory controller directly accesses the DIMMs on the same processor/memory board at the DIMM speed. The aggregate bandwidth for locally accessed memory increases with the number of processors.

The Opteron chipset supports dual-channel memory, which reduces memory latency by increasing the bus width from 64 bits + check bit to128 bits + check bits. Each memory controller is 64 bits wide and is operated in lock-step to support the 128 bit interface. Because the ProLiant DL585 has dual-width memory channels, DIMMs must be installed in pairs.

Overall memory speed and capacity for the DL585 is a function of the electrical loading and timing constraints of the various types of memory. There are two different processor/memory boards used in the ProLiant DL585 models. An 8-DIMM board supports PC2100 and PC2700 memory. A 4-DIMM memory board supports PC3200 memory. The chart below breaks out the different configuration options:

Table 2. Memory Configuration Options in the ProLiant DL585 server

Memory type	Max capacity	DIMM size	Maximum DIMMs per processor	Memory speed	Processor speeds
PC2100	64 GB	512 MB 1 GB 2 GB	8	266 MHz	1.6 – 2.4 GHz single-core
PC2700	128 GB	512 MB 1 GB 2 GB 4 GB	8	266 MHz	≥ 2.4 GHz single-core, all dual-core
PC2700	48 GB	512 MB 1 GB 2 GB	6	333 MHz	≥ 2.4 GHz single-core, all dual-core
PC2700	64 GB	4 GB	4	333 MHz	≥ 2.4 GHz single-core, all dual-core
PC3200	32 GB	512 MB 1 GB 2 GB	4	400 MHz	≥ 2.6 GHz single-core, all dual-core

To provide optimum performance for a wide variety of applications, the ProLiant DL585 can support either of two methods of organizing memory access: linear, non-uniform memory access (NUMA), or node interleaving sufficiently uniform memory accessing (SUMA).

A node consists of the processor, including the embedded memory controller and the attached DIMMs. The total memory attached to all the processors is divided into 4096 byte segments. In the case of linear addressing (NUMA), consecutive 4096 byte segments are on the same node. In the case of node interleaving (SUMA), consecutive 4096 byte segments are on different or adjacent nodes.

- Linear memory accessing (NUMA) defines the memory starting at 0 on node 0 and assigns the total amount of memory on node 0 the next sequential address, up to the memory total on node 0. The memory on node 1 will then start with the next sequential address until the process is complete.

- Node interleaving (SUMA) breaks memory into 4-KB addressable entities. Addressing starts with address 0 on node 0 and sequentially assigns through address 4095 to node 0, addresses 4096 through 8191 to node 1, addresses 8192 through 12287 to node 3, and addresses 12888 through 16383 to node 4. Address 16384 is assigned to node 0 and the process continues until all memory has been assigned in this fashion.

There is no rule for organizing memory for the best performance of an application because the difference in latency between the NUMA operation and SUMA operation is small. In general, a NUMA-aware operating system such as Microsoft Windows, and a NUMA-aware application such as Microsoft SQL Server will benefit from the NUMA organization. A NUMA-aware operating system and applications which allocate and de-allocate memory at the thread level will benefit from the NUMA organization, because the allocation and the thread will have a tendency to run on the same node. If an application uses a common allocation thread it will benefit from node interleaving.

The DL585 has, by default, NUMA memory configuration. For those applications that cannot take advantage of the NUMA architecture, performance may be improved by activating node interleaving (SUMA). System administrators can activate node interleaving using the HP ROM-Based Setup Utility (RBSU) provided as part of the ProLiant Essentials Foundation Pack.

I/O subsystem

I/O devices connect to the processor/memory boards via AMD's HyperTransport technology links rather than traditional I/O buses. AMD HyperTransport technology links are capable of signaling with clock speeds of up to 1 GHz, and Double Data Rate (DDR) memory signaling. Two factors contribute significantly to the signal speed:

- A HyperTransport point-to-point link, which speeds up data transfer by interconnecting all processors and embedded memory in the system. HyperTransport technology provides a universal connection designed to reduce the number of buses within the system, to provide a high-performance link for embedded applications, and to enable highly scalable multiprocessing systems. In the ProLiant DL585, HyperTransport technology delivers 8 GB/s processor-to-processor throughput for maximum performance and scalability.
- Resources do not share an I/O bus, so there is no overhead for bus arbitration. HyperTransport links can provide an effective throughput of 2.0 gigatransfers per pin-pair on a link.

The ProLiant DL585 system board has an integrated 1280 x 1024, 16M color graphics controller with 8 MB of SDRAM video memory on the PCI local bus.

The ProLiant DL585 server contains eight full-length, 64-bit PCI-X slots: Slots 1 and 2 operate at up to 133 MHz; slots 3 through 8 operate at up to 100 MHz. None of these slots supports PCI Hot-Plug Technology.

NOTE
The two PCI-X slots running at 133 MHz are limited to adapters with three PCI functions or less.

Figure 2. ProLiant DL585 PCI-X slots

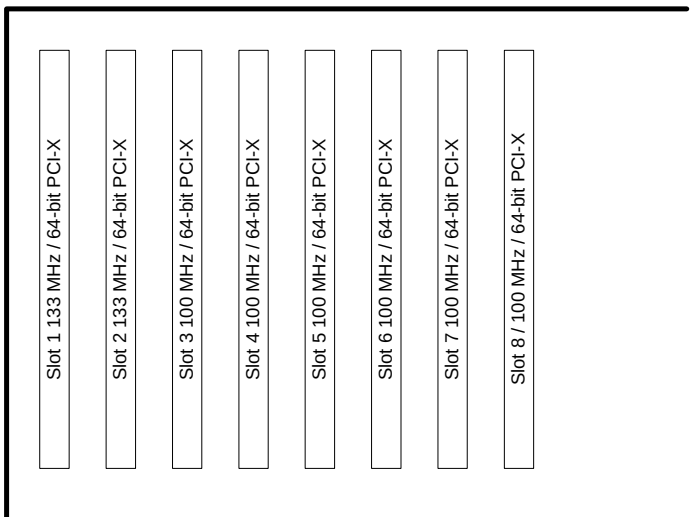

Performance balancing is the paired arrangement of expansion boards for optimal performance based on the bus architecture of the expansion slots. Properly balancing the boards across buses can improve performance. To balance expansion boards, populate slots across different buses before populating two slots on the same bus. The table below shows the preferred order of installation.

Table 3. Populating PCI-X expansion slots

Preferred installation order	Slot	Bus	Description
First	Slot 1	Fifth bus	133 MHz
Second	Slot 2	Sixth bus	133 MHz
Third	Slot 7	Third bus	100 MHz
Fourth	Slot 3	Seventh bus	100 MHz
Fifth	Slot 5	Eighth bus	100 MHz
Sixth	Slot 8	Third bus	100 MHz
Seventh	Slot 4	Seventh bus	100 MHz
Eighth	Slot 6	Eighth bus	100 MHz

Networking

The ProLiant DL585 server has an integrated, dual-port NC7782 Gigabit Server Adapter. The NC7782 is a PCI-X based LOM (LAN on Motherboard) embedded chip that supports 10/100/1000 Mbps Ethernet speeds and a PCI-X 64-bit/100-MHz data path that is compatible with existing PCI bus architectures. This LOM supports the standard HP network controller drivers and functionality, including teaming and failover to the other embedded NIC port or to stand-up PCI-X expansion card-based NICs.

NIC teaming support

The NC7782 supports two types of teaming: Network Fault Tolerance (NFT) teaming and Load Balancing teaming.

NFT teaming increases system availability by providing failover between a primary and a secondary network port on a server. It also monitors the server connection to the network and automatically switches traffic to a redundant link in the event of a failure. In a system configured for failover, a secondary port can take over for a failed primary port with no interruption of server operations.

Load balancing teaming, sometimes called port-bonding or trunking, enhances server availability and performance by sharing traffic among multiple network adapters. Load balancing capability allows optimum throughput by "bonding" multiple ports in a server to act as a single port. Load balancing teaming increases network bandwidth and redundancy.

High availability technologies

The ProLiant DL585 server supports dual, hot-pluggable, redundant power supplies to ensure uninterrupted service. All models ship standard with dual 870-Watt Hot-Plug Redundant Power Supplies (featuring hi/lo line auto-sensing).

The ProLiant DL585 supports internal hot-plug drives and has a transportable Battery-Backed Write Cache (BBWC) Enabler module with battery charge of up to 72 hours. This BBWC module protects up to 64 MB of write cache memory from hard boot, power, SCSI controller, or system board failures. In the event of a failure, write cache data can be moved to another ProLiant DL585 server in the data center by simply removing the BBWC Enabler and 5i Plus Memory Module simultaneously.

CAUTION

To prevent data loss, the cable between the BBWC Enabler and 5i Plus Memory Module must NOT be disconnected.

The ProLiant DL585 has eight hot-pluggable fans. The 7+1 redundant fan configuration provides a single cooling zone across the major server subsystems for7x24 cooling.

Storage technologies

The ProLiant DL585 server supports internal hot-plug storage and has an empty Wide Ultra3/Ultra320 drive cage. An internal Wide Ultra3/Ultra 320 backplane supports up to four 1-inch Ultra320 SCSI hot-plug drives and is configurable to support simplex (4x1) or duplex (2x2) mode. Duplex is the default mode. The server supports serial attached SCSI (SAS), small form factor (SFF) drives within the 8-Bay SAS/SATA Cage Option Kit. The SAS/SATA Cage must be used in conjunction with a SAS/SATA Controller such as the Smart Array P600 Controller to upgrade all U320 models to SAS. The server supports up to 1.20 TB of internal storage using 4 x 300 GB Ultra320 SCSI 1-inch hard drives or up to 582 GB of internal storage using 8 x 72.8 GB SAS SFF drives. The server supports Ultra320 drives with the addition of a PCI-X or PCI option card cabled to the backplane. An external SCSI port is not available on the ProLiant DL585 server. External SCSI solutions require the addition of a PCI or PCI-X option card with an external connector.

A dual-channel, Ultra3 Smart Array 5i Plus Controller is integrated on the DL585 system board. This controller supports RAID 0, 1, 1+0, and 5 across the internal hard disk drives. The DL585 supports a BBWC Enabler module on the Smart Array 5i Plus Controller with 64 MB of total

memory for code, transfer buffers, and read/write cache.[3] Optional SCSI cable assembly kit 365483-B21 is required for connecting a Smart Array Controller in a PCI slot to the internal hard drive backplane. Duplex configuration requires two such kits.

Management technologies

The ProLiant DL585 server is a highly reliable machine designed for efficient remote management and control. It ships standard with a full complement of HP management tools. It supports management tools including HP Integrated Lights-Out (iLO) management via an ASIC on the system board and ProLiant Essentials Foundation Pack.

HP Integrated Lights-Out

The server contains an HP Integrated Lights-Out (iLO) management ASIC on the system board and ships with iLO Standard Management functionality. That functionality is upgradeable through the purchase of an Integrated Lights-Out Advanced Pack. Table 3 summarizes the two levels of iLO management functionality.

Details about iLO management functionality are available at this URL: www.hp.com/servers/lights-out

Table 4. Functionality supported by HP Integrated Lights-Out management ASIC in the ProLiant DL585 server

iLO Standard Management	iLO Advanced Management
Virtual Text Remote Console	Virtual Graphical Remote Console
Virtual Power Button Control	USB-Based Virtual Media
Dedicated LAN connectivity	Directory Integration (Firmware Version 1.40 or greater)
Automatic IP configuration via DHCP/DNS/WINS	Terminal Services Integration (Firmware Version 1.5 or greater)
Industry-standard 128-bit Secure Sockets Layer (SSL) security	
IML and iLO event logging	
Support for 12 user accounts with customizable access privileges	

ProLiant Essentials Foundation Pack Software

The ProLiant Essentials Foundation Pack is a full suite of management software products. For more details, please see the following URL: www.hp.com/servers/proliantessentials.

Mechanical technologies

The ProLiant DL585 server leverages chassis and design of the ProLiant DL580 G2 server. The space-saving 4U form factor provides space efficiencies while providing maximum deployment flexibility. It is the first 4U server to support up to 4 AMD Opteron processors and up to 32 DIMMs.

[3] For a complete list of devices supported by the Smart Array 5i Controller see:
http://h18000.www1.hp.com/products/quickspecs/10890_div/10890_div.HTML (Worldwide)
http://h18000.www1.hp.com/products/quickspecs/10890_na/10890_na.HTML (North America)

Serviceability

The DL585 has a blue Unit ID light on the front and rear panels to assist technicians in locating a specific server within a rack. A front panel LED displays internal and external server health information. A QuickFind Diagnostic Display viewable through the top access panel of the server schematically represents potential sources of error for all major subsystems of the server (PCI-X I/O, memory, processor, redundant cooling, interlock, and platform over-temperature condition) and gives instant visual indication of fault status.

The modularity and virtually cable-less interior design of the server provide easy, tool-free access to all major subsystems and components. The server features front access to the power supplies. Removal of the system board is tool free. PCI card guides lock full-length PCI cards in place for transit and snap open for accessing the cards. The SCSI duplex/simplex drive cage connects directly to system board without cables. Technicians can configure this drive cage for simplex or duplex operation by means of simple switch. Eight individually removable fans connect without cables to the system board.

Conclusion

The ProLiant DL585 server is a 4U rack-optimized, 4-way server created for large data center deployments requiring enterprise-class performance, uptime, and scalability, plus ease of management and expansion. It offers customers running 32-bit applications increased performance and memory addressability. It provides a path to more powerful, 64-bit computing while protecting x86 investments.

For more information

For additional information, refer to the resources listed below.

Type of information	Source
ProLiant DL585 server QuickSpecs	http://h18004.www1.hp.com/products/quickspecs/11902_div/11902_div.html
Explanation of how HyperTransport technology works and the advantages if offers for system design and performance	AMD white paper *HyperTransport™ Technology: Simplifying System Design*, October 2002 www.hypertransport.org/docs/26635A_HT_System_Design.pdf
The AMD processor roadmap for HP ProLiant servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00428708/c00428708.pdf
Up-to-date information on operating systems and versions supported by the ProLiant DL585 server	http://www.hp.com/go/supportos
Latest drivers available for the HP ProLiant DL585 server	www.hp.com/support/files
ProLiant DL585 server documentation	http://h20000.www2.hp.com/bizsupport/TechSupport/Home.jsp?&lang=en&cc=us&prodTypeId=15351&prodSeriesId=398220&lang=en&cc=us
Technology brief on characterizing x86 processors for industry-standard servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00238028/c00238028.pdf
Details about HP Systems Insight Manager	www.hp.com/go/hpsim
Details about iLO management functionality	www.hp.com/servers/lights-out
Details about HP ProLiant Essentials	www.hp.com/servers/proliantessentials

Call to action

To help us better meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

TC060204TB, February 2006